An Iteratively Reweighted Norm Algorithm for Minimization of Total Variation Functionals

Brendt Wohlberg* and Paul Rodríguez

EDICS: IMD-ANAL

***Abstract*— Total Variation (TV) regularization has become a popular method for a wide variety of image restoration problems, including denoising and deconvolution. A number of authors have recently noted the advantages of replacing the standard ℓ^2 data fidelity term with an ℓ^1 norm. We propose a simple but very flexible method for solving a generalized TV functional which includes both the ℓ^2-TV and and ℓ^1-TV problems as special cases. This method offers competitive computational performance for ℓ^2-TV, and is comparable to or faster than any other ℓ^1-TV algorithms of which we are aware.**

***Index Terms*— image restoration, inverse problem, regularization, total variation**

I. Introduction

Total Variation (TV) regularization has become a very popular method for a wide variety of image restoration problems, including denoising and deconvolution [1], [2]. The standard ℓ^2-TV regularized solution of the inverse problem involving data $\mathbf{s}$ and forward linear operator K (the identity in the case of denoising, and a convolution for a deconvolution problem, for example) is the minimum of the functional

$$T(\mathbf{u}) = \frac{1}{p}\left\|K\mathbf{u} - \mathbf{s}\right\|_p^p + \frac{\lambda}{q}\left\|\sqrt{(D_x\mathbf{u})^2 + (D_y\mathbf{u})^2}\right\|_q^q \quad (1)$$

for $p = 2$ and $q = 1$, where we employ the following notation:

- the p-norm of vector $\mathbf{u}$ is denoted by $\|\mathbf{u}\|_p$,
- scalar operations applied to a vector are considered to be applied element-wise, so that, for example, $\mathbf{u} = \mathbf{v}^2 \Rightarrow u_k = v_k^2$ and $\mathbf{u} = \sqrt{\mathbf{v}} \Rightarrow u_k = \sqrt{v_k}$, and
- horizontal and vertical discrete derivative operators are denoted by D_x and D_y respectively.

While a number of algorithms [1], [3] have been proposed to solve this optimization problem, it remains a computationally expensive task which can be prohibitively costly for large problems and forward operators K without a fast implicit implementation or a sparse explicit matrix representation.

Recently, the ℓ^1-TV functional [4], [5], corresponding to Equation (1) with $p = 1$ and $q = 1$, has attracted attention due to a number of advantages [6], including superior performance with non-Gaussian noise such as speckle noise, and applications in cDNA microarray image processing [7], and cartoon-texture decomposition [8]. While a number of efficient algorithms have been proposed for the ℓ^2-TV problem (see, for example [1]), effective methods for the ℓ^1-TV problem have only recently been developed, the fastest of which [9], [10], [11] being applicable only to the $K = I$ denoising problem.

We propose a simple but computationally efficient and very flexible method for solving the generalized TV functional of Equation (1) for $p \geq 1$, $q \geq 1$, and K a linear operator which is not necessarily the identity.

II. Iteratively Reweighted Norm Approach

Our Iteratively Reweighted Norm (IRN) approach is motivated by the Iteratively Reweighted Least Squares (IRLS) algorithm [12], [13], [14] for solving the minimum ℓ^p norm problem

$$\min_{\mathbf{u}} \frac{1}{p}\left\|K\mathbf{u} - \mathbf{s}\right\|_p^p \quad (2)$$

by solving a sequence of minimum weighted ℓ^2 norm problems, and is also closely related to the Iterative Weighted Norm Minimization algorithms [15], [16] for sparse signal decompositions.

These methods represent the ℓ^p norm of $\mathbf{u}$

$$\frac{1}{p}\left\|\mathbf{u}\right\|_p^p = \frac{1}{p}\sum_k |u_k|^p,$$

by the weighted ℓ^2 norm of $\mathbf{u}$

$$\frac{1}{2}\left\|W^{1/2}\mathbf{u}\right\|_2^2 = \frac{1}{2}\mathbf{u}^T W \mathbf{u} = \frac{1}{2}\sum_k w_k u_k^2$$

with diagonal weight matrix $W = (2/p)\,\mathrm{diag}\left(|\mathbf{u}|^{p-2}\right)$. At each iteration of an iterative scheme, the ℓ^p norm is approximated by the weighted ℓ^2 norm using the weights from the previous iteration. To simplify somewhat, this approximation may be used to minimize the norm because, for the same choice of W (and $\mathbf{u}$ such that $u_k \neq 0\ \forall k$) we have

$$\nabla_{\mathbf{u}}\frac{1}{p}\left\|\mathbf{u}\right\|_p^p = (p/2)\nabla_{\mathbf{u}}\frac{1}{2}\left\|W^{1/2}\mathbf{u}\right\|_2^2,$$

so that both expressions have the same value and tangent direction. The proof of convergence for IRN is too long to reproduce here but is very similar to that for IRLS [14], and will be presented in detail in a paper that we are currently preparing.

Brendt Wohlberg is with T-7 Mathematical Modeling and Analysis, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. Email: brendt@t7.lanl.gov, Tel: (505) 667 6886, Fax: (505) 665 5757

Paul Rodríguez is with T-7 Mathematical Modeling and Analysis, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. Email: prodrig@t7.lanl.gov, Tel: (505) 606 1483 , Fax: (505) 665 5757

This work was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396 and was partially supported by the NNSA's Laboratory Directed Research and Development Program.

A. Data Fidelity Term

The data fidelity term of Equation (1) has the form of the IRLS functional in Equation (2), and is handled in the same way, representing

$$\frac{1}{p}\left\|K\mathbf{u}-\mathbf{s}\right\|_p^p \quad \text{by} \quad \frac{1}{2}\left\|W_F^{1/2}\left(K\mathbf{u}-\mathbf{s}\right)\right\|_2^2$$

with iteratively updated weights W_F. To avoid infinite weights for $p < 2$ and $u_k = 0$, we set

$$W_F = \operatorname{diag}\left(\frac{2}{p} f_F(K\mathbf{u}-\mathbf{s})\right)$$

where

$$f_F(x) = \begin{cases} |x|^{p-2} & \text{if } |x| > \epsilon_F \\ \epsilon_F^{p-2} & \text{if } |x| \leq \epsilon_F, \end{cases}$$

for some small ϵ_F, a common approach for IRLS algorithms [13]. In the limit as $\epsilon_F \to 0$, this weighted ℓ^2 norm tends to the original ℓ^p norm fidelity term.

B. Regularization Term

It is not quite as obvious how to express the regularization term from Equation (1) as a weighted ℓ^2 norm. Given vectors $\mathbf{u}$ and $\mathbf{v}$ we have (using block-matrix notation)

$$\left\|\begin{pmatrix} W^{1/2} & 0 \\ 0 & W^{1/2} \end{pmatrix}\begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}\right\|_2^2 = \sum_k w_k u_k^2 + w_k v_k^2$$

so that when

$$W = \operatorname{diag}\left(\frac{2}{q}\left(\mathbf{u}^2+\mathbf{v}^2\right)^{(q-2)/2}\right)$$

we have

$$\frac{1}{2}\left\|\begin{pmatrix} W^{1/2} & 0 \\ 0 & W^{1/2} \end{pmatrix}\begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}\right\|_2^2 = \frac{1}{q}\left\|\sqrt{\mathbf{u}^2+\mathbf{v}^2}\right\|_q^q .$$

We therefore define the operator D and weights $\tilde{W}$

$$D = \begin{pmatrix} D_x \\ D_y \end{pmatrix} \quad \tilde{W} = \begin{pmatrix} W & 0 \\ 0 & W \end{pmatrix}$$

so that $\|\tilde{W_R}^{1/2} D\mathbf{u}\|_2^2 = \|W_R^{1/2} D_x\mathbf{u}\|_2^2 + \|W_R^{1/2} D_y\mathbf{u}\|_2^2$ with weights defined by

$$W_R = \operatorname{diag}\left(\frac{2}{q}\left((D_x\mathbf{u})^2+(D_y\mathbf{u})^2\right)^{(q-2)/2}\right)$$

gives the desired term. (This is *not* the anisotropic separable approximation $\|D_x\mathbf{u}\|_q^q + \|D_y\mathbf{u}\|_q^q$ to $\left\|\sqrt{(D_x\mathbf{u})^2+(D_y\mathbf{u})^2}\right\|_q^q$ that is often used, e.g. [17].)

As in the case of the data fidelity term, care needs to be taken when $q < 2$ and $u_k = 0$. We define

$$f_R(x) = \begin{cases} |x|^{(q-2)/2} & \text{if } |x| > \epsilon_R \\ 0 & \text{if } |x| \leq \epsilon_R, \end{cases}$$

for some small ϵ_R and set

$$W_R = \operatorname{diag}\left(\frac{2}{q} f_R\left((D_x\mathbf{u})^2+(D_y\mathbf{v})^2\right)\right).$$

In the limit as $\epsilon_R \to 0$, this weighted ℓ^2 norm tends to the original ℓ^q norm fidelity term. Note that f_R sets values smaller than the threshold, ϵ_R, to zero, as opposed to f_F, which sets values smaller than the threshold, ϵ_F, to ϵ_F^{p-2}. Our motivation for this choice is that a region with very small or zero gradient should be allowed to have zero contribution to the regularization term, rather than be clamped to some minimum value. In practice, however, we have found that this choice does not give significantly different results than the standard IRLS approach represented by f_F.

C. General Algorithm

Combining the terms described in Sections II-A and II-B, we have the functional

$$T(\mathbf{u}) = \frac{1}{2}\left\|W_F^{1/2}\left(K\mathbf{u}-\mathbf{s}\right)\right\|_2^2 + \frac{\lambda}{2}\left\|\tilde{W_R}^{1/2} D\mathbf{u}\right\|_2^2$$

which, it is worth noting, may be expressed as

$$T(\mathbf{u}) = \frac{1}{2}\left\|\begin{pmatrix} W_F^{1/2} & 0 \\ 0 & \tilde{W_R}^{1/2} \end{pmatrix}\left(\begin{pmatrix} K \\ \sqrt{\lambda}D \end{pmatrix}\mathbf{u} - \begin{pmatrix} \mathbf{s} \\ 0 \end{pmatrix}\right)\right\|,$$

which has the same form as an IRLS problem, but differs in the computation of $\tilde{W_R}^{1/2}$. The minimum of this functional is

$$\mathbf{u} = \left(K^T W_F K + \lambda D^T \tilde{W}_R D\right)^{-1} K^T W_F \mathbf{s}, \qquad (3)$$

and the resulting algorithm consists of the following steps:

Initialize

$$\mathbf{u}_0 = \left(K^T K + \lambda D^T D\right)^{-1} K^T \mathbf{s}$$

Iterate

$$\begin{aligned} W_{F,k} &= \operatorname{diag}\left(\frac{2}{p} f_F(K\mathbf{u}_{k-1}-\mathbf{s})\right) \\ W_{R,k} &= \operatorname{diag}\left(\frac{2}{q} f_R\left((D_x\mathbf{u}_{k-1})^2+(D_y\mathbf{u}_{k-1})^2\right)\right) \\ \mathbf{u}_k &= \left(K^T W_{F,k} K + \lambda D_x^T W_{R,k} D_x \right. \\ &\quad \left. + \lambda D_y^T W_{R,k} D_y\right)^{-1} K^T W_{F,k}\mathbf{s} \end{aligned}$$

The matrix inversion is achieved using the Conjugate Gradient (CG) method. We have found that a significant speed improvement may be achieved by starting with a high CG tolerance which is decreased with each main iteration until the final desired value is reached.

D. Denoising Algorithm

In the case of the denoising problem, when $K = I$, we may apply the substitution $\tilde{\mathbf{u}} = W_F^{1/2}\mathbf{u}$, giving

$$T(\tilde{\mathbf{u}}) = \frac{1}{2}\left\|\tilde{\mathbf{u}} - W_F^{1/2}\mathbf{s}\right\|_2^2 + \frac{\lambda}{2}\left\|\tilde{W_R}^{1/2} D W_F^{-1/2}\tilde{\mathbf{u}}\right\|_2^2,$$

with solution

$$\tilde{\mathbf{u}} = \left(I + \lambda W_F^{-1/2} D^T \tilde{W}_R D W_F^{-1/2}\right)^{-1} W_F^{1/2}\mathbf{s}. \qquad (4)$$

Applying this modification to the general algorithm of Section II-C was found to result (see Section III) in a very large reduction in the required number of CG iterations.

III. Computational Results

In the remainder of this paper we shall restrict our attention to the ℓ^1-TV case ($p = 1$, $q = 1$), but note that this flexible approach is capable of efficiently solving other cases as well, including the standard ℓ^2-TV case ($p = 2$, $q = 1$) where we have found it to be slightly slower than the lagged diffusivity algorithm [1], to which it is related. For the results reported here, operators D_x and D_y were defined by applying the same one-dimensional discrete derivative along image rows and columns respectively. Applied to vector $\mathbf{u} \in \mathbb{R}^N$, this discrete derivative was computed as $u_k - u_{k+1}$ for $k \in \{0, 1, \ldots, N-2\}$, and set to zero at $N - 1$ (equivalent to a half-sample symmetric boundary extension). Constants ϵ_F and ϵ_R were set to values in the range 10^{-1} to 10^{-3}, representing a reasonable compromise between solution quality (the weighted ℓ^2 approximations to the ℓ^p and ℓ^q norms deteriorate as these values become larger) and run time (the linear system in Equation (3) becomes increasingly poorly conditioned as they become smaller). Except where specified otherwise, run times were obtained on a 3GHz Intel Pentium 4 processor.

A. Denoising

We tested the denoising performance of the IRN algorithm on the 512×512 pixel Lena image with 10% of the pixels corrupted by speckle noise, giving an SNR of 1.2dB. Figures 1 and 2 display the reconstruction qualities (with respect to the reference noise-free image) and run times respectively for three different choices of regularization parameter λ. Note that, for the best choice of $\lambda = 1.25$, the SNR curve has flattened out by iteration 6, corresponding to a run time of 4.5s. (The curve of functional values against iteration number, omitted due to space constraints, is also almost flat by iteration 6.) The advantages of the denoising-specific algorithm (Section II-D) over the general algorithm (Section II-C) are clearly shown in Figure 3.

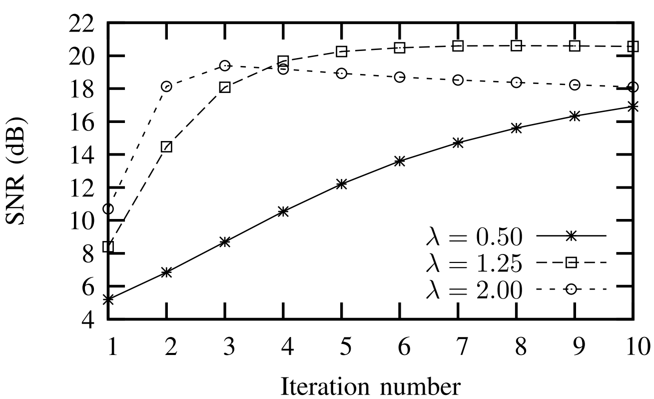

Fig. 1. Denoising SNR values against algorithm iteration number.

A run-time performance comparison with the Goldfarb and Yin [18] method is provided in Table I, the parameters of both methods having been selected for a similar quality/run-time tradeoff. (This comparison is for a downsampled Lena image since the SOCP solver exhibited convergence problems for the full 512×512 image used elsewhere in this paper.) Note that

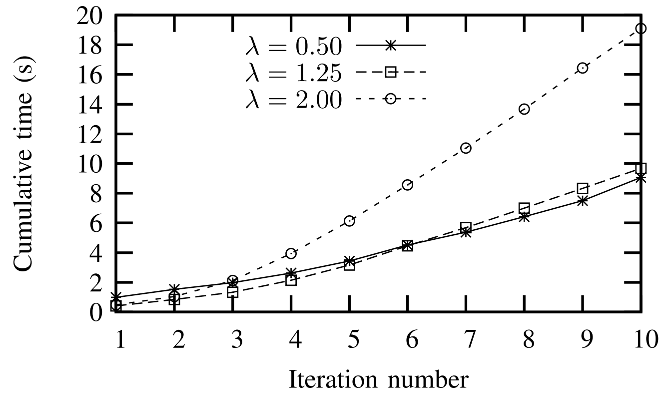

Fig. 2. Denoising time against algorithm iteration number.

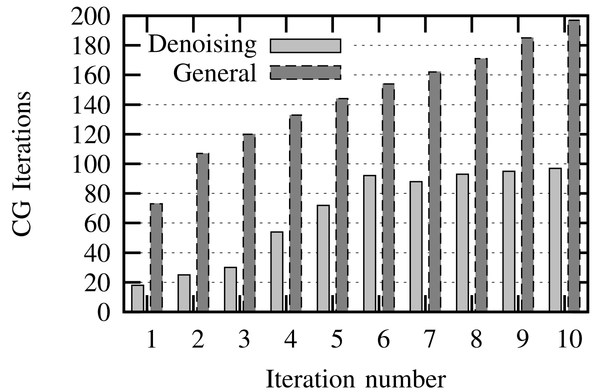

Fig. 3. A comparison of CG iterations for the denoising (Section II-D) and general (Section II-C) algorithms.

the IRN method is more than an order of magnitude faster, for this application. We have not been able to conduct direct performance comparisons with the recent Darbon and Sigelle method [10], [11], but comparisons between their published run times and our own experiments performed on comparable platforms suggest that their method is of comparable speed to the IRN method. Their method has the advantage of providing an exact solution to the TV optimization problem, but this is usually not an issue for practical denoising application, and furthermore, is not extensible to more general image restoration problems.

B. Deconvolution

We use a deconvolution problem to demonstrate the application of the IRN algorithm when $K \neq I$, choosing K as the

Method	Time (s)
SOCP [18]	20.4
IRN (general)	2.1
IRN (denoising)	0.5

TABLE I
A comparison between SOCP and IRN ℓ^1-TV denoising run times for a 256×256 Lena image corrupted by 5% speckle noise. Results were computed on a 2.2GHz AMD Opteron 148 CPU.

linear operator corresponding to convolution by a separable smoothing filter having 9 taps and approximating a Gaussian with standard deviation of 2.0. In this case, since $K \neq I$, the substitution applied in the previous section is no longer possible. We used the same test image as in Section III-A, convolving it with the filter described above, and corrupted 10% of the pixels with speckle noise, giving a 1.4dB SNR image. Figure 4 displays the reconstruction SNR values against algorithm iteration number for three different values of λ, while Figure 5 display the evolution of the TV functional and its weighted approximation for the best choice $\lambda = 10^{-3}$. Note that both of these curves have flattened out by iteration 7, which corresponds to a run time of 209s. (Using a variable CG solver accuracy strategy, we are able to achieve a similar SNR in 118s.) We are not aware of any other algorithms with comparable performance for general ($K \neq I$) reconstruction problems.

IV. Conclusions

The IRN approach provides a simple but computationally efficient method for TV regularized optimization problems, including both denoising and those having a linear operator in the data fidelity term, such as deconvolution. This method is very flexible, and can be applied to regularized inversions with a wide variety of norms for the data fidelity and regularization terms, including the standard ℓ^2 TV, and more recently proposed ℓ^1 TV formulations, and, in particular, provides a very fast algorithm for the ℓ^1 TV case. A software implementation [19] is available under an open-source license.

Acknowledgment

The authors thank Wotao Yin for his generous assistance in preparing the SOCP method run time comparisons, and Markus Berndt for valuable discussion on CG preconditioning methods.

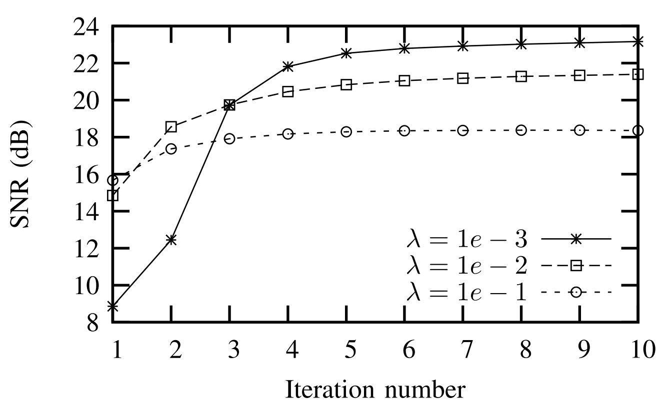

Fig. 4. Deconvolution SNR values against algorithm iteration number.

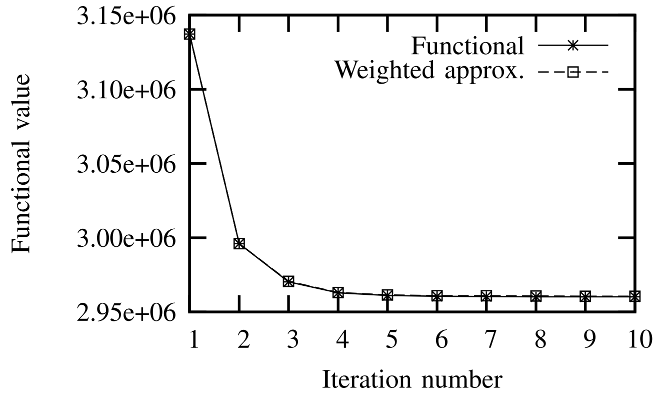

Fig. 5. Comparison of functional values and weighted ℓ^2 approximations for deconvolution with $\lambda = 10^{-3}$. (Curves are indistinguishable at this scale).

References

[1] C. R. Vogel and M. E. Oman, "Iterative methods for total variation denoising," *SIAM J. on Scientific Computing*, vol. 17, no. 1-4, pp. 227–238, Jan. 1996.

[2] T. Chan, S. Esedoglu, F. Park, and A. Yip, "Recent developments in total variation image restoration," in *The Handbook of Mathematical Models in Computer Vision*, N. Paragios, Y. Chen, and O. Faugeras., Eds. Springer, 2005.

[3] A. Chambolle, "An algorithm for total variation minimization and applications," *J. of Math. Imaging and Vision*, vol. 20, pp. 89–97, 2004.

[4] S. Alliney, "Digital filters as absolute norm regularizers," *IEEE Trans. on Signal Processing*, vol. 40, no. 6, pp. 1548–1562, 1992.

[5] M. Nikolova, "Minimizers of cost-functions involving nonsmooth data-fidelity terms. application to the processing of outliers," *SIAM J. on Numerical Analysis*, vol. 40, no. 3, pp. 965–994, 2002.

[6] T. F. Chan and S. Esedoglu, "Aspects of total variation regularized L^1 function approximation," *SIAM J. on Applied Mathematics*, vol. 65, no. 5, pp. 1817–1837, 2005.

[7] W. Yin, T. Chen, X. S. Zhou, and A. Chakraborty, "Background correction for cDNA microarray images using the TV+L^1 model," *Bioinformatics*, vol. 21, no. 10, pp. 2410–2416, 2005.

[8] W. Yin, D. Goldfarb, and S. Osher, "Image cartoon-texture decomposition and feature selection using the total variation regularized L^1 functional," *Lecture Notes in Comp. Sci.*, vol. 3752, pp. 73–84, 2005.

[9] J. F. Aujol, G. Gilboa, T. Chan, and S. Osher, "Structure-texture image decomposition - modeling, algorithms, and parameter selection," *Int. J. of Computer Vision*, vol. 67, no. 1, pp. 111–136, 2006.

[10] J. Darbon and M. Sigelle, "Image restoration with discrete constrained total variation part I: Fast and exact optimization," *J. of Math. Imaging and Vision*, vol. 26, no. 3, pp. 261–276, 2006.

[11] ——, "Image restoration with discrete constrained total variation part II: Levelable functions, convex priors and non-convex cases," *J. of Math. Imaging and Vision*, vol. 26, no. 3, pp. 277–291, 2006.

[12] A. E. Beaton and J. W. Tukey, "The fitting of power series, meaning polynomials illustrated on band-spectroscopic data," *Technometrics*, no. 16, pp. 147–185, 1974.

[13] J. A. Scales and A. Gersztenkorn, "Robust methods in inverse theory," *Inverse Problems*, vol. 4, no. 4, pp. 1071–1091, Oct. 1988.

[14] K. P. Bube and R. T. Langan, "Hybrid ℓ^1/ℓ^2 minimization with applications to tomography," *Geophysics*, vol. 62, no. 4, pp. 1183–1195, July-August 1997.

[15] I. F. Gorodnitsky and B. D. Rao, "A new iterative weighted norm minimization algorithm and its applications," in *IEEE Sixth SP Workshop on Statistical Signal and Array Processing, 1992*, Victoria, BC, Canada, Oct. 1992.

[16] ——, "Sparse signal reconstruction from limited data using FOCUSS: a re-weighted minimum norm algorithm." *IEEE Trans. on Signal Processing*, vol. 45, no. 3, pp. 600–616, Mar. 1997.

[17] Y. Li and F. Santosa, "A computational algorithm for minimizing total variation in image restoration," *IEEE Trans. on Image Processing*, vol. 5, no. 6, pp. 987–995, June 1996.

[18] D. Goldfarb and W. Yin, "Second-order cone programming methods for total variation based image restoration," *SIAM J. on Scientific Computing*, vol. 27, no. 2, pp. 622–645, 2005.

[19] P. Rodríguez and B. Wohlberg, "Numerical methods for inverse problems and adaptive decomposition (NUMIPAD)," Software library available from http://numipad.sourceforge.net/.